

A Stochastic Probing Problem with Applications

Anupam Gupta*

Viswanath Nagarajan†

Abstract

We study a general *stochastic probing* problem defined on a universe V , where each element $e \in V$ is “active” independently with probability p_e . Elements have weights $\{w_e : e \in V\}$ and the goal is to maximize the weight of a chosen subset S of active elements. However, we are given only the p_e values—to determine whether or not an element e is active, our algorithm must *probe* e . If element e is probed and happens to be active, then e must irrevocably be added to the chosen set S ; if e is not active then it is not included in S . Moreover, the following conditions must hold in every random instantiation:

- the set Q of probed elements satisfy an “outer” packing constraint,
- the set S of chosen elements satisfy an “inner” packing constraint.

The kinds of packing constraints we consider are intersections of matroids and knapsacks. Our results provide a simple and unified view of results in stochastic matching [12, 3] and Bayesian mechanism design [9], and can also handle more general constraints. As an application, we obtain the first polynomial-time $\Omega(1/k)$ -approximate “Sequential Posted Price Mechanism” under k -matroid intersection feasibility constraints, improving on prior work [9, 25, 19].

1 Introduction

We study an adaptive stochastic optimization problem along the lines of [20, 14, 13, 16]. The *stochastic probing* problem is defined on a universe V of elements with weights $\{w_e : e \in V\}$. We are also given two downwards-closed set systems (V, \mathcal{I}_{in}) and (V, \mathcal{I}_{out}) , which we call the *inner* and *outer* packing constraints, whose meanings we shall give shortly. For each element $e \in V$, there is a probability p_e , where element e is *active/present* with this probability, independently of all other elements. We want to choose a set $S \subseteq V$ of active elements belonging to \mathcal{I}_{in} , i.e., all elements in the chosen set S must be active and also independent according to the inner packing constraint ($S \in \mathcal{I}_{in}$). The goal is to maximize the expected weight of the chosen set.

However, the information about which elements are active and which are inactive is not given up-front. All we know are the probabilities p_e , and that the active set is a draw from the product distribution given by $\{p_e\}_{e \in V}$ —to determine if an element e is active or not, we must *probe* e . Moreover, if we probe e , and e happens to be active, then we *must* irrevocably add e to our chosen set S —we do not have a right to discard any probed element that turns out to be active. This “query and commit” model is quite natural in a number of applications such as kidney exchange, online dating and auction design (see below for details).

Finally, there is a constraint on which elements we can probe: the set Q of elements probed in any run of the algorithm must be independent according to the outer packing constraint \mathcal{I}_{out} —i.e., $Q \in \mathcal{I}_{out}$. This is the constraint that gives the probing problem its richness. Since every probed element that

*Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Supported in part by NSF award CCF-1016799 and an Alfred P. Sloan Fellowship.

†IBM T.J. Watson Research Center.

is active must be included in the solution which needs to maintain independence in \mathcal{I}_{in} , at any point t (with current solution S_t and currently probed set Q_t) we can only probe those elements e with $Q_t \cup \{e\} \in \mathcal{I}_{out}$ and $S_t \cup \{e\} \in \mathcal{I}_{in}$. (Indeed, if $p_e = 0$ there is no point probing e ; and if $p_e > 0$ there is a danger that e is active and we will be forced to add it to S_t , which we cannot if $S_t \cup \{e\} \notin \mathcal{I}_{in}$.)

While the stochastic probing problem seems fairly abstract, it has interesting applications: we give two applications of this problem, to designing posted-price Bayesian auctions, and to modeling problems in online dating/kidney exchange. We first state our results and then describe these applications.

1.1 Our Results

For the *unweighted* stochastic probing problem (i.e., $w_e = 1$ for all $e \in V$), if both inner and outer packing constraints are given by k -systems¹, we consider the greedy algorithm which considers elements in decreasing order of their probability p_e , probing them whenever feasible.

Theorem 1.1 (Unweighted Probing) *The greedy algorithm for unweighted stochastic probing achieves a tight $\frac{1}{k_{in}+k_{out}}$ -approximation ratio, when \mathcal{I}_{in} is a k_{in} -system and \mathcal{I}_{out} is a k_{out} -system.*

This result generalizes the greedy 4-approximation algorithm for unweighted stochastic matching, by Chen et al. [12], where both inner and outer constraints are b -matchings (and hence 2-systems). For the special case of stochastic matching, Adamczyk [1] gave an improved factor-2 bound. However, Theorem 1.1 is tight in our setting of general k -systems; its proof is LP-based, and we feel it is much simpler than previous proofs for the special cases. The main idea of our proof is a dual-fitting argument that extends the Fisher et al. [15] analysis of the greedy algorithm for k -matroid intersection. In Section 5 we generalize our unweighted probing result to also handle “global time” constraints, at the loss of a small constant factor in the approximation ratio (see Theorem 5.1).

There is no known greedy algorithm for stochastic probing in the *weighted* case (as opposed to the deterministic setting of finding the maximum weight set subject to a k -system, where greedy gives a $1/k$ -approximation [18, 15]); indeed, natural greedy approaches can be arbitrarily bad even for weighted stochastic matching [12]. Hence, we use an LP relaxation for the weighted probing problem, where variables correspond to marginal probabilities of probing/choosing elements in the optimal policy. This is similar to previous works on such adaptive stochastic problems [14, 13, 3]. Our rounding algorithm is based on the recently introduced notion of *contention resolution (CR) schemes* for packing constraints, due to Chekuri et al. [11]. Loosely speaking, given a packing constraint on a universe V and a fractional solution $\{x_e\}_{e \in V}$, a CR-scheme is a two-step rounding procedure where

- a. Each element e is chosen independently into $I_1 \subseteq V$ with probability proportional to x_e .
- b. A feasible subset $I_2 \subseteq I_1$ (suitably computed) is output as the solution.

We show that the existence of suitable CR-schemes for both \mathcal{I}_{in} and \mathcal{I}_{out} imply an approximation algorithm for weighted stochastic probing, where the approximation ratio depends on the quality of the two CR-schemes. Our main result for weighted stochastic probing is Theorem 3.4 (which requires some notation to state precisely), but here is a representative corollary:

Theorem 1.2 (Weighted Probing: Special Case) *There is an $\Omega\left(\frac{1}{k_{in}+k_{out}}\right)$ -approximation algorithm for weighted stochastic probing when the inner and outer constraints are intersections of k_{in}*

¹For any integer k , a k -system is a downwards-closed collection of sets $\mathcal{I} \subseteq 2^V$ such that for any $S \subseteq V$, the maximal subsets of S that belong to \mathcal{I} can differ in size by at most a factor of k . Examples are intersections of k matroids, and k -set packing.

and k_{out} matroids, respectively. Moreover, there is an $\Omega\left(\frac{1}{(k_{in}+k_{out})^2}\right)$ -approximation algorithm under arbitrary k_{in} and k_{out} system constraints.

Some of the other allowed constraints are unsplittable flow on trees (under the “no-bottleneck” assumption) and packing integer programs. Details on the weighted case appear in Section 3.

1.2 Applications

We now give two applications: the first shows how our algorithm for the weighted probing problem immediately gives us posted price auctions for single parameter settings where the feasibility set is given by intersections of matroids, the second is an application for dating/kidney exchange. Both of these extend and generalize previous results in these areas.

Bayesian Auction Design. Consider a mechanism design setting for a single seller facing n single-parameter buyers. The seller has a feasibility constraint given by a downward-closed set system $\mathcal{I} \subseteq 2^{[n]}$ and is allowed to serve any set of buyers from \mathcal{I} . Buyers are *single-parameter*; i.e., buyer i ’s private data is a single real number v_i which denotes his valuation of being served (if i is not served then he receives zero value). In the *Bayesian* setting, the valuation v_i is drawn from some set $\{0, 1, \dots, B\}$ according to probability distribution \mathcal{D}_i ; here we assume that the valuations of buyers are discrete and independently drawn. The valuation v_i is private to the buyer, but the distribution \mathcal{D}_i is public knowledge. The goal in these problems is a revenue-maximizing truthful mechanism that accepts bids from buyers and outputs a feasible allocation (i.e., a set $S \in \mathcal{I}$ of buyers that receive service), along with a price that each buyer has to pay for service. A very special type of mechanism is a *Sequential Posted Pricing Mechanism* (SPM) that chooses a price for each buyer and makes “take-it-or-leave-it” offers to the buyers in some order [22, 8, 9]. Such mechanisms are simple to run and obviously truthful (see [9] for a discussion of other advantages), hence it is of interest to design SPMs which achieve revenue comparable to the revenue-optimal mechanism.

Designing the best SPM can be cast as a stochastic probing problem on a universe $V = \{1, 2, \dots, n\} \times \{0, 1, \dots, B\}$, where element (i, c) corresponds to offering a price c to buyer i . Element (i, c) has weight $w_{ic} = c$, which is the revenue obtained if the offer “price c for buyer i ” is accepted, and has probability $p_{ic} = \Pr_{v_i \sim \mathcal{D}_i}[v_i \geq c]$, which is the probability that i will indeed accept service at price c . The inner constraint \mathcal{I}_{in} is now the natural lifting of the actual constraints \mathcal{I} to the universe V , where $\{(i, c)\}_{c \geq 0}$ are copies of i . The outer constraint \mathcal{I}_{out} requires that at most one of the elements $\{(i, c) \mid c \geq 0\}$ can be probed for each i : i.e., each buyer i can be offered at most one price. This serves two purposes: firstly, it gives us a posted-price mechanism. Secondly, we required in our model that each element (i, c) is active with probability p_{ic} , *independently* of the other elements (i, c') ; however, the underlying semantics imply that if i accepts price c , then she would also accept any $c' \leq c$, which would give us correlations. Constraining ourselves to probe at most one element corresponding to each buyer i means we never probe two correlated elements, and hence the issue of correlations never arises.

Our results for stochastic probing give near-optimal SPMs for many feasibility constraints. Moreover, we show that our LP relaxation not only captures the best possible SPMs, but also captures the optimal truthful mechanism *of any form* under the Bayes-Nash equilibrium (and hence Myerson’s optimal mechanism [21]). In the case of k matroid intersection feasibility constraints, our results give the first polynomial-time sequential posted price mechanisms whose revenue is $\Omega(1/k)$ times the optimum. Previous papers [9, 25, 19] proved the existence of such SPMs, but they were polynomial-time only for $k \leq 2$. For larger k , previous works only showed *existence* of $\Omega(1/k)$ -approximate SPMs, and polynomial-time implementations of these SPMs only obtained an $\Omega(1/k^2)$ fraction of the optimal

revenue. The previous results compare the performance of their SPMs directly to the revenue of the optimal mechanism [21], whereas we compare our SPMs to an LP relaxation of this mechanism, which is potentially larger. Moreover, our general framework gives us more power:

- We can handle broader classes of feasibility constraints \mathcal{I} , not just matroid intersections: e.g., we can model auctions involving unsplittable flow on trees, which can be used to capture allocations of point-to-point bandwidths in a tree-shaped network. This is because the feasibility constraints \mathcal{I} for the auction directly translate into inner constraints for the probing problem.
- We can also handle additional side-constraints to the auction via a richer class of outer constraints \mathcal{I}_{out} . For example, the seller may incur costs in the form of time/money to make offers. Such budget limits can be modeled in the stochastic probing problem as an extra outer knapsack constraint, and our algorithm finds approximately optimal SPMs even in this case. More generally, our algorithm can easily handle a rich class of other resource constraints (matroid intersections, packing IPs etc) on the auction. However, in the presence of these side-constraints, our algorithm’s revenue is an approximation only to the best SPM satisfying these constraints, and no longer comparable to the unconstrained optimal mechanism.

Online dating and Kidney Exchange [12] Consider a dating agency with several users. Based on the profiles of users, the agency can compute the probability that any pair of users will be compatible. Whether or not a pair is successfully matched is only known after their date; moreover, in the case of a match, both users immediately leave the site (happily). Furthermore, each user has a patience/timeout level, which is the maximum number of failed dates after which he/she drops out of the site (unhappily). The objective of the dating site is to schedule dates so as to maximize the expected number of matched pairs. (Similar constraints arise in kidney exchange systems.) This can be modeled as stochastic probing with the universe V being edges of the complete graph whose nodes correspond to users. The inner constraints specify that the chosen edges be a matching in G . The outer constraints specify that for each node j , at most t_j edges incident to j can be probed, where t_j denotes the patience level of user j . Both these are b -matching constraints; in fact when the graph is bipartite, they are intersections of two partition matroids.

Our results will give an alternate way to obtain constant factor approximation algorithms for this stochastic matching problem. Such algorithms were previously given by [12, 3], but they relied heavily on the underlying graph structure. Additionally, our techniques allow for more general sets of constraints. E.g., not all potential dates may be equally convenient to a user, and (s)he might prefer dates with other nearby users. This can be modeled as a sequence of patience bounds for the user, specifying the maximum number of dates that the user is willing to go outside her neighborhood/city/state etc. In particular, if u_1, u_2, \dots, u_n denote the users in decreasing distance from user j then there is a non-decreasing sequence $\langle t_j^1, \dots, t_j^n \rangle$ of numbers where user j wishes to date at most t_j^r users among the r farthest other users $\{u_1, \dots, u_r\}$. This corresponds to the stochastic probing problem, where the inner constraint remains matching but the outer constraint becomes a 2-system. Our algorithm achieves a constant approximation even here.

1.3 Other Related Work

Dean et al. [14, 13] were the first to consider approximation algorithms for *stochastic packing* problems in the adaptive optimization model. For the stochastic knapsack problem, where items have random sizes (that instantiate immediately after selection), [14] gave a $(3 + \epsilon)$ -approximation algorithm; this was improved to $2 + \epsilon$ in [6, 5]. [13] considered stochastic packing integer programs (PIPs) and gave approximation guarantees matching the best known deterministic bounds. Our stochastic probing

problem can be viewed as a two-level generalization of stochastic packing, with two different packing constraints: one for probed elements, and one for chosen elements. However, all random variables in our setting are $\{0, 1\}$ -valued (each element is either active or not), whereas [14, 13] allow arbitrary non-negative random variables.

Chen et al. [12] first studied a *stochastic probing* problem: they introduced the unweighted stochastic matching problem and showed that greedy is a 4-approximation algorithm. Adamczyk [1] improved the analysis to show a bound of 2. Both these proofs involve intricate arguments on the optimal decision tree. In contrast, our analysis of greedy is much simpler and LP-based, and extends to the more general setting of k -systems. (For the stochastic matching, our result implies a 4-approximation.) Bansal et al. [3] gave a different LP proof that greedy is a 5-approximation for stochastic matching, but their proof relied heavily on the graph structure, making the extension to general k -systems unclear. [3] also gave the first $O(1)$ -approximation for *weighted* stochastic matching, which was LP-based. ([12] showed that natural greedy approaches for weighted stochastic matching are arbitrarily bad.) Our algorithm for weighted probing is also LP-based, where we make use of the elegant abstraction of “contention resolution schemes” introduced by Chekuri et al. [11] (see Section 3), which provides a clean approach to rounding the LP.

The papers of Chawla et al. [9], Yan [25], and Kleinberg and Weinberg [19] study the performance of Sequential Posted Price Mechanisms (SPMs) for Bayesian single-parameter auctions, and relate the revenue obtained by SPMs to the optimal (non-posted-price) mechanism given by Myerson [21]. Our algorithm for stochastic probing also yields SPMs for Bayesian auctions where the feasible sets of buyers are specified by, e.g., k -matroid intersection and unsplittable flow on trees. Our proof relates an LP relaxation of the optimal mechanism to the LP used for stochastic probing. Linear programs have been used to model optimal auctions in a number of settings; e.g., see Vohra [24]. Bhattacharya et al. [7] also used LP relaxations to obtain approximately optimal mechanisms in a Bayesian setting with multiple items and budget constrained buyers.

1.4 Preliminaries

Specifying Probing Algorithms. A solution (policy) to the stochastic probing problem is an *adaptive* strategy of probing elements satisfying the constraints imposed by \mathcal{I}_{out} and \mathcal{I}_{in} . At any time step $t \geq 1$, let Q_t denote the set of elements already probed and S_t the current solution (initially $Q_1 = S_1 = \emptyset$); an element $e \in V \setminus Q_t$ can be probed at time t if and only if $Q_t \cup \{e\} \in \mathcal{I}_{out}$ and $S_t \cup \{e\} \in \mathcal{I}_{in}$. If e is probed then exactly one of the following happens:

- e is active (with probability p_e), and $Q_{t+1} \leftarrow Q_t \cup \{e\}$, $S_{t+1} \leftarrow S_t \cup \{e\}$, or
- e is inactive (with probability $1 - p_e$), and $Q_{t+1} \leftarrow Q_t \cup \{e\}$, $S_{t+1} \leftarrow S_t$.

Hence the policy is a decision tree with nodes representing elements that are probed and branches corresponding to their random instantiations. Note that an optimal policy may be exponential sized, and designing a polynomial-time algorithm requires tackling the question of whether there exist polynomial-sized near-optimal strategies. A *non adaptive* policy is simply given by a permutation on V , where elements are considered in this order and probed whenever feasible in both \mathcal{I}_{out} and \mathcal{I}_{in} . The *adaptivity gap* compares the best non-adaptive policy to the best adaptive policy.

Packing Constraints. We model packing constraints as *independence systems*, which are of the form $(V, \mathcal{I} \subseteq 2^V)$ where V is the *universe* and \mathcal{I} is a collection of *independent sets*. We assume \mathcal{I} is *downwards closed*, i.e., $A \in \mathcal{I}$ and $B \subseteq A \implies B \in \mathcal{I}$. Some examples are:

- *Knapsack constraint*: each element $e \in V$ has size $s_e \in [0, 1]$ and $\mathcal{I} = \{A \subseteq V \mid \sum_{e \in A} s_e \leq 1\}$.
- *Matroid constraint*: an independence system (V, \mathcal{I}) where for any subset $S \subseteq V$, every maximal

independent subset of S has the same size. See [23] for many properties and examples.

- *k-system*: an independence system (V, \mathcal{I}) where for any subset $S \subseteq V$, every *maximal* independent subset of S has size at least $\frac{1}{k}$ times the size of the *maximum* independent subset of S . For example: matroids are 1-systems, matchings are 2-systems, and intersections of k matroids form k -systems.
- *Unsplittable Flow Problem (UFP) on trees*: there is an edge-capacitated tree T , and each element $e \in V$ corresponds to a path P_e in T and demand d_e . Subset $S \subseteq V$ is independent (i.e. $S \in \mathcal{I}$) iff $\{\text{path } P_e \text{ with demand } d_e\}_{e \in S}$ is routable in T . We assume the “no-bottleneck” condition, where the maximum demand $\max_{e \in V} d_e$ is at most the minimum capacity in T .

When the universe is clear from context, we refer to an independence system (V, \mathcal{I}) just as \mathcal{I} . We also make use of linear programming relaxations for independence systems: the LP relaxation of \mathcal{I} is denoted by $\mathcal{P}(\mathcal{I}) \subseteq [0, 1]^V$ and contains the convex hull of all independent sets. (Since $\mathcal{P}(\mathcal{I})$ is a relaxation it need not equal the convex hull). For example: $\mathcal{P}(\mathcal{I}) = \{\mathbf{x} \in [0, 1]^V : \sum_{e \in V} s_e \cdot x_e \leq 1\}$ for knapsacks; $\mathcal{P}(\mathcal{I}) = \{\mathbf{x} \in [0, 1]^V : \sum_{e \in S} x_e \leq r_{\mathcal{I}}(S), \forall S \subseteq V\}$ for matroids, where $r_{\mathcal{I}}(\cdot)$ denotes the rank function.

1.5 Outline

We first consider the unweighted probing problem in Section 2. Then, in Section 3 we study the weighted probing problem. In Section 4 we present the application to posted price mechanisms for Bayesian auctions (Theorem 4.1). Finally, in Section 5 we study the generalization of unweighted probing to the setting of global time constraints.

2 Unweighted Stochastic Probing

In this section, we study the stochastic probing problem with unit weights, i.e., $w_e = 1$ for all $e \in V$. We assume the inner and outer packing constraints are a k_{in} -system and a k_{out} -system, respectively. We show that the greedy algorithm, which considers elements in non-increasing order of their probabilities p_e and probes them when feasible, has performance claimed in Theorem 1.1. We give an LP-based dual-fitting proof of this result.

For brevity, let us use k to denote k_{in} , and k' to denote k_{out} . Let the *rank function* of \mathcal{I}_{in} be $r : 2^V \rightarrow \mathbb{N}$, where for each $S \subseteq V$, $r(S) = \max\{|I| \mid I \in \mathcal{I}, I \subseteq S\}$ be the *maximum* size of an independent subset of S . By definition of k -systems, for any $S \subseteq V$, any maximal independent set of S (according to \mathcal{I}_{in}) has size at least $r(S)/k$. Similarly, let $r' : 2^V \rightarrow \mathbb{N}$ denote the rank function of \mathcal{I}_{out} . We may not be able to evaluate the rank function, since this is NP-complete for $k \geq 3$. For any $T \subseteq V$, let $\text{span}(T) = \{e \in V : r(T \cup \{e\}) = r(T)\}$ be the *span* of T . Likewise, let span' denote the span function for \mathcal{I}_{out} .

Claim 2.1 *For any $T \subseteq V$, the maximum independent subset of T (which has size $r(T)$) is a maximal independent subset of $\text{span}(T)$. Hence, for $T \subseteq V$ and $R \subseteq V$, we have $r(\text{span}(T)) \leq k \cdot r(T) \leq k \cdot |T|$ and $r'(\text{span}'(R)) \leq k' \cdot r'(R) \leq k' \cdot |R|$.*

Let us write the natural LP relaxation and dual for the probing problem:

$$\begin{array}{ll} \max & \sum_{e \in V} p_e y_e \\ \text{s.t.} & \sum_{e \in S} p_e y_e \leq r(S) \quad \forall S \subseteq V \\ & \sum_{e \in S} y_e \leq r'(S) \quad \forall S \subseteq V \\ & \mathbf{y} \geq 0. \end{array} \quad \left| \quad \begin{array}{ll} \min & \sum_S r(S) \alpha(S) + \sum_S r'(S) \beta(S) \\ \text{s.t.} & p_e \sum_{S: e \in S} \alpha(S) + \sum_{S: e \in S} \beta(S) \geq p_e \quad \forall e \in V \\ & \alpha(S), \beta(S) \geq 0 \quad \forall S \subseteq V. \end{array} \right.$$

Claim 3.1 in the next section shows that this LP is a valid relaxation. It is not known if these linear programs can be solved in polynomial time for arbitrary p -systems \mathcal{I}_{in} and \mathcal{I}_{out} ; we use them only

for the analysis. Note that the greedy algorithm defines a non-adaptive strategy. Consider a sample path π down the natural decision tree associated with the above algorithm; it is completely defined by the randomness in which elements are active. Let $\Pr[\pi]$ denote its probability, and Q_π, S_π be the sets probed and picked on taking this path.

Lemma 2.2 *If alg is the random variable denoting the number of elements picked,*

$$\mathbb{E}[\text{alg}] = \sum_{\pi} \Pr(\pi) \cdot |S_\pi| = \sum_{\pi} \Pr(\pi) \cdot \sum_{e \in Q_\pi} p_e.$$

Proof: The first equality follows by definition of expectations, and the fact that elements are unweighted. For the second, let $\pi_{<e}$ be the outcomes of elements before e in the ordering. Note that the event $\mathbf{1}(e \text{ probed})$ is completely determined by $\pi_{<e}$. Moreover,

$$\Pr[e \text{ picked} \mid \pi_{<e}] = \mathbf{1}(e \text{ probed} \mid \pi_{<e}) \cdot p_e.$$

Hence, the expected value of the algorithm is

$$\begin{aligned} \mathbb{E}[\text{alg}] &= \sum_e \sum_{\pi_{<e}} \Pr[\pi_{<e}] \cdot \Pr[e \text{ picked} \mid \pi_{<e}] = \sum_e \sum_{\pi_{<e}} \Pr[\pi_{<e}] \cdot \mathbf{1}(e \text{ probed} \mid \pi_{<e}) \cdot p_e \\ &= \sum_e \sum_{\pi} \Pr[\pi] \cdot \mathbf{1}(e \text{ probed} \mid \pi_{<e}) \cdot p_e = \sum_e \sum_{\pi} \Pr[\pi] \cdot \sum_{e \in Q_\pi} p_e. \end{aligned}$$

Above, we used the fact that e 's being probed (or equivalently, it's lying in Q_π) was purely a function of $\pi_{<e}$. And that e being active is independent of all others. \blacksquare

Lemma 2.3 *For each outcome π , there is a feasible dual of value at most $k|S_\pi| + k' \sum_{e \in Q_\pi} p_e$. Moreover, there is a feasible dual of value at least $(k + k')\mathbb{E}[\text{alg}]$.*

The following proof is similar to that of Fisher et al. [15] showing that the greedy algorithm is a k -approximation for the intersection of k matroids.

Proof: Let $A = \text{span}(S_\pi)$ be the span of the set of picked elements S_π ; note that by Claim 2.1, $r(A) \leq k \cdot |S_\pi|$. We set $\alpha(A) = 1$, and all other α variables to zero.

Let the set of probed elements $Q_\pi = \{a_1, a_2, \dots, a_\ell\}$ in this order. Define

$$\beta(\text{span}'(\{a_1, a_2, \dots, a_h\})) := p_{a_h} - p_{a_{h+1}} \geq 0$$

for all $h \in \{1, \dots, \ell\}$ (where we imagine $p_{a_{\ell+1}} = 0$). This is also well-defined since every subset of Q_π is independent in \mathcal{I}_{out} . The non-negativity follows from the greedy algorithm that probes elements in decreasing probabilities. The dual objective value equals:

$$r(A) + \sum_{h=1}^{\ell} r'(\text{span}'(\{a_1, a_2, \dots, a_h\})) \cdot (p_{a_h} - p_{a_{h+1}}) \leq k \cdot |S_\pi| + \sum_{h=1}^{\ell} k' \cdot h \cdot (p_{a_h} - p_{a_{h+1}}),$$

which is $k \cdot |S_\pi| + k' \sum_{e \in Q_\pi} p_e$. The inequality is by Claim 2.1. Next we show that the dual solution is feasible. The non-negativity is clearly satisfied, so it remains to check feasibility of the dual covering constraints. For any $e \in V$,

- Case I: $e \in Q_\pi$. Say $e = a_g$ in the ordering of the set Q_π . Then e lies in $\text{span}'(\{a_1, a_2, \dots, a_h\})$ for all $h \geq g$. Hence, the left hand side of e 's covering constraint contributes at least

$$\sum_{h=g}^{\ell} \beta(\text{span}'(\{a_1, a_2, \dots, a_h\})) = \sum_{h=g}^{\ell} (p_{a_h} - p_{a_{h+1}}) = p_{a_g} = p_e.$$

- Case II: $e \notin Q_\pi$ because of the outer constraint. Say e was seen when the Q set was $\{a_1, a_2, \dots, a_g\}$. Then $e \in \text{span}'(\{a_1, a_2, \dots, a_h\})$ for all $h \geq g$. In this case, the left hand side contributes at least

$$\sum_{h=g}^{\ell} \beta(\text{span}'(\{a_1, a_2, \dots, a_h\})) = \sum_{h=g}^{\ell} (p_{a_h} - p_{a_{h+1}}) = p_{a_g} \geq p_e.$$

Here we used the fact that elements are considered in decreasing order of their probabilities.

- Case III: $e \notin Q_\pi$ because of the inner constraint. Then $e \in \text{span}(S_\pi) = A$, and hence the $p_e \sum_{S:e \in S} \alpha(S) = p_e \alpha(A) = p_e$.

This proves the first part of the lemma. Taking expectations over π , the resulting convex combination $\sum_{\pi} \Pr[\pi](\alpha_\pi, \beta_\pi)$ of these feasible duals is another feasible dual of value $k \mathbb{E}[|S_\pi|] + k' \mathbb{E}[\sum_{e \in Q_\pi} p_e]$, which by Lemma 2.2 equals $(k + k') \mathbb{E}[\text{alg}]$. ■

Our analysis for the greedy algorithm is tight. In particular, if all p_e 's equal one, and the inner and outer constraints are intersections of (arbitrary) partition matroids, then we obtain the greedy algorithm for $(k_{in} + k_{out})$ -dimensional matching. The approximation ratio in this case is known to be exactly $k_{in} + k_{out}$.

Application to Unweighted Stochastic Matching. When the inner constraint is matching (which is a 2-system) and the outer constraint is b -matching (also a 2-system) on the same graph, we obtain the unweighted stochastic matching problem of Chen et al. [12]. Hence Theorem 1.1 gives an alternate proof of greedy being a 4-approximation [12]. We know now that greedy is a 2-approximation [1], but we currently do not know an LP-based proof of this bound.

3 Weighted Stochastic Probing

We now turn to the general weighted case of stochastic probing. Here the natural combinatorial algorithms perform poorly, so we use linear programming relaxations of the problem, which we round to get non-adaptive policies. Given an instance of the stochastic probing problem with inner constraints (V, \mathcal{I}_{in}) and outer constraints (V, \mathcal{I}_{out}) , we use the following LP relaxation:

$$\begin{aligned} \max \quad & \sum_{e \in V} w_e \cdot x_e \\ \text{s.t.} \quad & x_e = p_e \cdot y_e \quad \forall e \in V \\ & x \in \mathcal{P}(\mathcal{I}_{in}) \\ & y \in \mathcal{P}(\mathcal{I}_{out}) \end{aligned} \quad (\mathcal{LP})$$

We assume that the LP relaxations of the inner and outer constraints can be solved efficiently: this is true for matroids, knapsacks, UFP on trees, and their intersections. For general k -systems, it is not known if this LP can be solved exactly. However, using the fact that the greedy algorithm achieves a $\frac{1}{k}$ -approximation for maximizing linear objective functions over k -systems (even with respect to the LP relaxation, which follows from [15], or the proof of Lemma 2.3), and the equivalence of approximate separation and optimization [17], we can obtain a $\frac{1}{\max\{k_{in}, k_{out}\}}$ -approximate LP solution when \mathcal{I}_{in} and \mathcal{I}_{out} are arbitrary k_{in} and k_{out} systems.

Claim 3.1 *The optimal value of $(\mathcal{LP}) \geq$ optimal value of the probing instance.*

Proof: Let y_e^* denote the probability that element e is probed by the optimal strategy; i.e., $y_e^* = \Pr[e \in Q^*]$. Also let x_e^* denote the probability that element e is chosen in the final solution, $x_e^* =$

$\Pr[e \in S^*]$. Due to the constraints, we have $Q^* \in \mathcal{I}_{out}$ and $S^* \in \mathcal{I}_{in}$, and hence $y^* \in \mathcal{P}(\mathcal{I}_{out})$ and $x^* \in \mathcal{P}(\mathcal{I}_{in})$. Moreover,

$$x_e^* = \Pr[e \in S^*] = \Pr[e \in Q^* \text{ and } e \text{ active}] = p_e \cdot \Pr[e \in Q^*] = p_e \cdot y_e^*, \quad \forall e \in V.$$

Here we used the fact that the probability of element e being active is independent of the past decisions, and in particular, of the optimal strategy's decision to probe e . Thus (x^*, y^*) is a feasible solution to \mathcal{LP} . Finally, the optimal value of the probing problem instance is $\sum_{e \in V} w_e \cdot \Pr[e \in S^*] = \sum_e w_e x_e^*$, which is the LP objective value of (x^*, y^*) . ■

3.1 Contention-Resolution Schemes

Given a solution (x, y) for the LP relaxation, we need to get a policy from it. Our rounding algorithm is based on the elegant abstraction of *contention resolution schemes (CR schemes)*, as defined in Chekuri et al. [11]. Here is the formal definition, and the main theorem we will use.

Definition 3.2 *An independence system $(V, \mathcal{J} \subseteq 2^V)$ with LP-relaxation $\mathcal{P}(\mathcal{J})$ admits a monotone (b, c) **CR-scheme** if, for any $z \in \mathcal{P}(\mathcal{J})$ there is a (possibly randomized) mapping $\pi : 2^V \rightarrow \mathcal{J}$ such that:*

- (i) *If $I \subseteq V$ is a random subset where each element $e \in V$ is chosen independently with probability $b \cdot x_e$, $\Pr_{I, \pi}[e \in \pi(I) \mid e \in I] \geq c$ for all $e \in V$.*
- (ii) *For any $e \in I_1 \subseteq I_2 \subseteq V$, $\Pr_\pi[e \in \pi(I_1)] \geq \Pr_\pi[e \in \pi(I_2)]$.*
- (iii) *The map π can be computed in polynomial time.*

Moreover, $\pi : 2^V \rightarrow \mathcal{J}$ is a (b, c) **ordered CR-scheme** if there is a (possibly random) permutation σ on V so that for each $I \subseteq V$, $\pi(I)$ is the maximal independent subset of I obtained by considering elements in the order of σ .

Theorem 3.3 ([11, 10, 4, 9]) *There are monotone CR-schemes for the following independence systems (below, $0 < b \leq 1$ is any value unless specified otherwise)*

- $(b, (1 - e^{-b})/b)$ CR-scheme for matroids.
- $(b, 1 - k \cdot b)$ ordered CR-scheme for k -systems.
- $(b, 1 - 6b)$ ordered CR-scheme for unsplittable flow on trees, with the “no bottleneck” assumption, for any $0 < b \leq 1/60$.
- $(b, 1 - 2kb)$ CR-scheme for k -column sparse packing integer programs.

The CR-scheme for k -systems can be inferred from Lemma 4.12 in [11] using the observation that $r(\text{span}(R)) \leq k \cdot |R|$ for any $R \subseteq V$ in a k -system.

3.2 How to Round the LP Solution

Given the formalism of CR schemes, we can now state our main result for rounding a solution to the relaxation (\mathcal{LP}) .

Theorem 3.4 *Consider any instance of the stochastic probing problem with*

- (i) (b, c_{out}) CR-scheme for $\mathcal{P}(\mathcal{I}_{out})$.
- (ii) Monotone (b, c_{in}) **ordered CR-scheme** for $\mathcal{P}(\mathcal{I}_{in})$.

Then there is a $b \cdot (c_{out} + c_{in} - 1)$ -approximation algorithm for the weighted stochastic probing problem.

Before we prove Theorem 3.4, we observe that combining Theorems 3.4 and 3.3 gives us, for example:

- a $1/(4(k + \ell))$ -approximation algorithm when the inner and outer constraints are intersections of k and ℓ matroids respectively.
- an $\Omega(1)$ -approximation algorithm when the inner and outer constraints are unsplittable flows on trees/paths satisfying the no-bottleneck assumption.
- an $\Omega(1/(k + \ell)^2)$ -approximation algorithm when the inner and outer constraints are arbitrary k and ℓ systems. Here, we lose an additional $k + \ell$ factor in solving \mathcal{LP} approximately.

The Rounding Algorithm. Let π_{out} denote the randomized mapping corresponding to a (b, c_{out}) CR-scheme for $y \in \mathcal{P}(\mathcal{I}_{out})$, and π_{in} be that corresponding to a (b, c_{in}) CR-scheme for $x \in \mathcal{P}(\mathcal{I}_{in})$. The algorithm to round the LP solution (x, y) for weighted stochastic probing appears as Algorithm 3.1.

Algorithm 3.1 Rounding Algorithm for Weighted Probing

- 1: Pick $I \subseteq 2^V$ by choosing each $e \in V$ independently with probability $b \cdot y_e$.
 - 2: Let $P = \pi_{out}(I)$. (By definition of the CR scheme, $P \in \mathcal{I}_{out}$ with probability one.)
 - 3: Order elements in P according to σ (the inner ordered CR scheme) to get $e_1, e_2, \dots, e_{|P|}$.
 - 4: Set $S \leftarrow \emptyset$.
 - 5: **for** $i = 1, \dots, |P|$ **do**
 - 6: **if** $(S \cup \{e_i\} \in \mathcal{I}_{in})$ **then**
 - 7: Probe e_i : set $S \leftarrow S \cup \{e_i\}$ if e_i is active, and $S \leftarrow S$ otherwise.
-

The Analysis. We now show that $\mathbb{E}[w(S)]$ is large compared to the LP value $\sum_e w_e x_e$. To begin, a few observations about this algorithm. Note that this is a randomized strategy, since there is randomness in the choice of I and maybe in the maps π_{out} and π_{in} . Also, by the CR scheme properties, the probed elements are in \mathcal{I}_{out} , and the chosen elements in \mathcal{I}_{in} . Finally, having chosen the set P to (potentially) probe, the elements actually probed in step 7 relies on the *ordered* CR scheme for the inner constraints. In Appendix A we show that some simpler rounding algorithms that work for stochastic matching do not apply in this more general setting.

Recall that $I \subseteq V$ is the random set where each element e is included independently with probability $b \cdot y_e$; also $P = \pi_{out}(I)$. Let $J \subseteq V$ be the set of active elements; i.e., each $e \in V$ is present in J independently with probability p_e . The set of chosen elements is now $S = \pi_{in}(P \cap J)$. The main lemma is now:

Lemma 3.5 *For any $e \in V$,*

$$\Pr_{I, \pi_{out}, J, \pi_{in}} [e \in \pi_{in}(\pi_{out}(I) \cap J)] \geq b \cdot (c_{out} + c_{in} - 1) \cdot x_e,$$

where b, c_{out}, c_{in} are parameters given by our CR-schemes.

Proof: Recall that $P = \pi_{out}(I)$, so we want to lower bound:

$$\begin{aligned} \Pr[e \in \pi_{in}(P \cap J)] &= \Pr[e \in \pi_{in}(P \cap J) \wedge e \in I \cap J \cap P] \\ &= \Pr[e \in I \cap J \cap P] - \Pr[e \notin \pi_{in}(P \cap J) \wedge e \in I \cap J \cap P] \\ &\geq b x_e \cdot c_{out} - \Pr[e \notin \pi_{in}(P \cap J) \wedge e \in I \cap J \cap P], \end{aligned} \tag{1}$$

where the inequality uses $\Pr[e \in I \cap J] = b y_e \cdot p_e = b x_e$ and $\Pr[e \in P = \pi_{out}(I) | e \in I \cap J] \geq c_{out}$ by Definition 3.2(i) applied to the outer CR scheme, since I is a random subset chosen according to $b \cdot y$ where $y \in \mathcal{P}(\mathcal{I}_{out})$.

We now upper bound $\Pr[e \notin \pi_{in}(P \cap J) \wedge e \in I \cap J \cap P]$ by $(1 - c_{in}) \cdot bx_e$ which combined with (1) would prove the lemma. Now, condition on any instantiation $I = I_1$, $P = \pi_{out}(I_1) = P_1 \subseteq I_1$ and $J = J_1$ such that $e \in I_1 \cap J_1 \cap P_1$. Then,

$$\Pr[e \notin \pi_{in}(P_1 \cap J_1)] \leq \Pr[e \notin \pi_{in}(I_1 \cap J_1)], \quad (2)$$

by Definition 3.2(ii) applied to the inner CR scheme (since $e \in P_1 \cap J_1 \subseteq I_1 \cap J_1$). Taking a linear combination of the inequalities in (2) with respective multipliers $\Pr[I = I_1, J = J_1, P = P_1]$ (where $e \in I_1 \cap J_1 \cap P_1$), we obtain

$$\begin{aligned} \Pr[e \notin \pi_{in}(P \cap J) \wedge e \in I \cap J \cap P] &\leq \Pr[e \notin \pi_{in}(I \cap J) \wedge e \in I \cap J \cap P] \\ &\leq \Pr[e \notin \pi_{in}(I \cap J) \wedge e \in I \cap J] \\ &= bx_e \cdot \Pr[e \notin \pi_{in}(I \cap J) | e \in I \cap J] \end{aligned}$$

where the equality uses $\Pr[e \in I \cap J] = by_e \cdot p_e = bx_e$. The last expression above is at most $bx_e(1 - c_{in})$ by Definition 3.2(i) applied to the inner CR scheme, since $I \cap J$ is a random subset chosen according to $b \cdot x$ where $x \in \mathcal{P}(\mathcal{I}_{in})$. This proves $\Pr[e \notin \pi_{in}(P \cap J) \wedge e \in I \cap J \cap P] \leq (1 - c_{in}) \cdot bx_e$ as desired. ■

Consequently, the expected weight of the chosen set S is

$$\mathbb{E} \left[\sum_{e \in S} w_e \right] = \sum_{e \in V} w_e \cdot \Pr[e \in \pi_{in}(P \cap J)] \geq b(c_{in} + c_{out} - 1) \cdot \sum_{e \in V} w_e \cdot x_e.$$

The inequality uses Lemma 3.5. This completes the proof of Theorem 3.4.

Remark: We note that our results also hold in a slightly more general model where the elements are not necessarily independent, but every set $T \in \mathcal{I}_{out}$ is mutually independent.²

- Observe that \mathcal{LP} is a valid relaxation for stochastic probing, even in this setting. The only change in the proof of Claim 3.1 is: if Q^* and S^* denote the sets of probed and chosen elements in an optimal policy then $\Pr[e \in S^*] = \Pr[e \in Q^* \text{ and } e \text{ active}] = \Pr[e \in Q^*] \cdot \Pr[e \text{ active} | e \in Q^*] = \Pr[e \in Q^*] \cdot p_e$, where the last equality uses the fact that at any point in the optimal policy when e is probed, “element e being active” is independent of the previously observed elements (which along with e is some set in \mathcal{I}_{out} and hence is mutually independent).
- Moreover, in Lemma 3.5, if we let $J \subseteq V$ denote the random subset where each element e is present independently with probability p_e and $J_a \subseteq V$ the set of active elements, then $\Pr[e \in \pi_{in}(\pi_{out}(I) \cap J_a)] = \Pr[e \in \pi_{in}(\pi_{out}(I) \cap J)]$. This is because, conditioning on any $I = I_1$ and $P = \pi_{out}(I_1) = P_1$, the distributions of $J_a \cap P_1$ and $J \cap P_1$ are identical (by mutual independence of P_1).

4 Bayesian Single Parameter Mechanism Design

In this section, we show how a Bayesian single-parameter auction problem can be modeled as a stochastic probing problem, yielding new posted-price mechanisms for such auctions.

Formally, we consider a *Bayesian* mechanism design problem with one seller and n *single-parameter* agents that bid for service. The term “single-parameter” means that each agent i ’s private information is represented by a single number v_i , which is the agent’s *valuation*. In the Bayesian setting, the valuation $v_i \in \{0, 1, \dots, B\}$ is drawn from an independent probability distribution \mathcal{D}_i .³ These

²A set $\{E_i\}_{i=1}^\ell$ of events is mutually independent if for any subset $L \subseteq [\ell]$ we have $\Pr[\bigwedge_{i \in L} E_i] = \prod_{i \in L} \Pr[E_i]$.

³We can also handle continuous distributions by approximating them via discrete distributions, at the loss of a small constant factor.

valuations are private knowledge, but the distributions \mathcal{D}_i are publicly known. Each agent $i \in [n]$ submits a *bid* $b_i \in \{0, 1, \dots, B\}$ representing his valuation. The seller has a *feasibility constraint* given by a downward closed set system $\mathcal{I} \subseteq 2^{[n]}$, and hence can serve any set of agents from \mathcal{I} .

A *mechanism* is a function that maps a bid-vector $\mathbf{b} \in \{0, 1, \dots, B\}^n$ to an *allocation* $A(\mathbf{b}) \in \mathcal{I}$, along with *prices* $\pi_i(\mathbf{b})$ to be paid by each agent $i \in [n]$. For notational convenience, we define $X_i(\mathbf{b}) := \mathbf{1}_{i \in A(\mathbf{b})}$ denoting whether or not agent i receives service. The *utility* of agent i under bids \mathbf{b} is $v_i \cdot X_i(\mathbf{b}) - \pi_i(\mathbf{b})$. We consider mechanisms satisfying the following standard properties:

- *Voluntary participation*: an agent pays only when receiving service and the payment is at most his bid. For all bid vectors \mathbf{b} , $\pi_i(\mathbf{b}) \leq X_i(\mathbf{b}) \cdot b_i$ for all agents $i \in [n]$.
- *No positive transfers*: the mechanism does not pay agents, $\pi_i(\mathbf{b}) \geq 0$ for all $i \in [n]$ and \mathbf{b} .
- *Truthful in expectation*: For each $i \in [n]$ and $v_i \in \{0, 1, \dots, B\}$, if v_i is agent i 's true valuation then his expected utility by bidding v_i is at least his expected utility under any other bid $b_i \in \{0, 1, \dots, B\}$, i.e.,

$$\mathbb{E}_{b_j \leftarrow \mathcal{D}_j: j \neq i} [v_i \cdot X_i(\mathbf{b}_{-i}, v_i) - \pi_i(\mathbf{b}_{-i}, v_i)] \geq \mathbb{E}_{b_j \leftarrow \mathcal{D}_j: j \neq i} [v_i \cdot X_i(\mathbf{b}_{-i}, b_i) - \pi_i(\mathbf{b}_{-i}, b_i)]$$

We are interested in designing a mechanism that maximizes the expected revenue. The well-known *Myerson mechanism* [21] is optimal for the single-parameter setting, and it proceeds by reducing the revenue maximization problem to the welfare-maximization setting (which can then be solved using the VCG mechanism). However the resulting mechanism can be complicated to implement, and is computationally hard under combinatorial feasibility constraints such as intersections of more than two matroids. Hence, simpler mechanisms such as “sequential posted price mechanisms” (SPM) are often desirable in practice; see [9, 25, 19] for further discussion on this. In an SPM, the seller offers “take it or leave it” prices to the agents one-by-one. When the feasible set \mathcal{I} is given by the intersection of k matroid constraints, Chawla et al. [9] showed the existence of SPMs achieving a $\frac{1}{k+1}$ -approximation to the optimal mechanism. However, when the posted prices are to be computed in polynomial time, the approximation ratio becomes $\Omega(1/k^2)$. Indeed, getting $\Omega(1/k)$ -approximate SPMs for intersections of $k \geq 2$ matroids was an open problem before this work.

In this section, we show that approximately optimal SPMs can be obtained as an application of the stochastic probing problem. Since our algorithm for computing prices runs in polynomial time, we obtain a polynomial time SPM for k -matroid constraints that is an $\Omega(1/k)$ -approximation to the optimal mechanism. We proceed by first showing that computing the optimal posted price mechanism is an instance of the matroid constrained probing problem. Then we show that the LP relaxation of this probing problem (which we use for our algorithm) has value at least that of the optimal (potentially non-posted price) mechanism as well, which completes the argument.

4.1 Posted Price Mechanism as Probing Problem

Given the distributions \mathcal{D}_i of each agent $i \in [n]$, and the feasibility constraint \mathcal{I} , we are interested in computing a *sequential posted price mechanism*. This corresponds to setting prices π_i for each agent $i \in [n]$, and making “take it or leave it” offers to agents in a suitable sequence (while ensuring feasibility in \mathcal{I}). A rational agent will accept an offer if and only if the posted price is at most his valuation. It is clear that any such mechanism is truthful since the prices are independent of bids. In fact this mechanism is truthful even when each agent knows the precise bids of all other agents, which is a stronger condition than truthfulness in expectation.

Consider an instance of the stochastic probing problem with:

- Universe $V := \{(i, c) : i \in [n], c \in \{0, 1, \dots, B\}\}$.
- Weights $w_{i,c} = c$ for all $(i, c) \in V$.
- Probabilities $p_{i,c} = \Pr_{v_i \leftarrow \mathcal{D}_i} [v_i \geq c]$ for all $(i, c) \in V$.
- The outer constraint being a partition matroid: \mathcal{I}_{out} consists of all subsets $S \subseteq V$ with $|S \cap \{(i, c)\}_{c=0}^B| \leq 1$ for all $i \in [n]$. This corresponds to offering at most one price to each agent.
- The inner constraint being the natural lifting of the seller's feasibility constraint (on universe $[n]$) to V , where $\{(i, c)\}_{c=0}^B$ are copies of i . Formally, \mathcal{I}_{in} consists of all subsets $S \subseteq V$ with (a) $|S \cap \{(i, c)\}_{c=0}^B| \leq 1$ for all $i \in [n]$ and (b) $\{i \in [n] : \exists (i, c) \in S\} \in \mathcal{I}$.

Notice that if \mathcal{I} is given by an intersection of p matroids then so is \mathcal{I}_{in} .

Due to the outer constraint, a solution to this probing problem never probes two copies of the same agent. This ensures two properties: (1) the independence assumption on elements of V agrees with the auction setting where copies $\{(i, c)\}_{c=0}^B$ of each agent i are actually dependent, and (2) we obtain a posted price mechanism. Moreover, the inner constraint handles the feasibility constraint \mathcal{I} . Thus, solutions to this probing problem correspond precisely to sequential posted price mechanisms and vice versa. Instead of modeling the elements $\{(i, c)\}_{c=0}^B$ of each agent i as being active independently, we could also model their joint distribution induced by \mathcal{D}_i : since every set in \mathcal{I}_{out} is now mutually independent, Theorem 3.4 still applies (as noted in the end of Section 3).

We can now use our algorithm for the weighted stochastic probing problem to obtain an approximately optimal SPM. In the next subsection, we show that the optimal revenue (of any mechanism, which may potentially be non-posted-price) is at most the value of the stochastic probing LP. Since our approximation ratio for stochastic probing is relative to this LP, we obtain the following result (setting $b = \frac{1}{2k+1}$, $c_{in} = 1 - kb$ and $c_{out} = (1 - e^{-b})/b \geq 1 - b/2$ in Theorem 3.4)

Theorem 4.1 *There is a polynomial-time sequential posted price mechanism for k matroid intersection constraints, which has revenue at least $\frac{1}{4k+2}$ times the revenue of the optimal mechanism.*

More generally, this result holds for any feasibility constraint \mathcal{I} that admits an *ordered* CR scheme, where the approximation ratio depends on the quality of the CR scheme. For example, this also implies a constant factor approximate SPM when \mathcal{I} is given by an unsplittable flow on trees.

4.2 Bounding the Optimal Mechanism

Recall that our algorithm for the weighted probing problem is based on the LP relaxation \mathcal{LP} . For instances corresponding to the Bayesian mechanism design problem (from the reduction above), this LP is:

$$LP_P = \max \sum_{i \in [n]} \sum_{c=0}^B c \cdot x_{i,c} \quad (3)$$

$$\text{subject to } x_{i,c} = \Pr[v_i \geq c] \cdot y_{i,c} \quad \forall (i, c) \in V \quad (4)$$

$$\left\{ \sum_{c=0}^B x_{i,c} : i \in [n] \right\} \in \mathcal{P}(\mathcal{I}) \quad (5)$$

$$\sum_{c=0}^B y_{i,c} \leq 1 \quad \forall i \in [n] \quad (6)$$

$$\mathbf{y} \geq \mathbf{0}. \quad (7)$$

Constraint (5) is the inner constraint which is a lifting of \mathcal{I} , and (6) is the outer partition matroid constraint. We will show that the optimal value of this LP is least the value of the optimal mechanism for the Bayesian auction problem. To do so, we want to write an LP relaxation for the optimal mechanism. Consider any mechanism given by allocations $\{X_i(\mathbf{b})\}_{i \in [n]}$ and prices $\{\pi_i(\mathbf{b})\}_{i \in [n]}$ as functions of bids. For each $i \in [n]$ and $c \in \{0, 1, \dots, B\}$, define:

$$z_{i,c} := \mathbb{E}_{b_j \leftarrow \mathcal{D}_j: j \neq i} [X_i(\mathbf{b}_{-i}, c)] \quad \text{and} \quad q_{i,c} := \mathbb{E}_{b_j \leftarrow \mathcal{D}_j: j \neq i} [\pi_i(\mathbf{b}_{-i}, c)].$$

For each agent i and value c , when agent i bids c , $z_{i,c}$ is the probability that i is served by the mechanism and $q_{i,c}$ is the expected price that i is charged (both expectations are taken over valuations of all other agents $[n] \setminus i$).

Lemma 4.2 (Myerson [21], Archer and Tardos [2]) *Any mechanism that satisfies truthfulness in expectation and voluntary participation has:*

- A. $z_{i,c}$ is non-decreasing in c , for all $i \in [n]$.
- B. $q_{i,c} \leq c \cdot z_{i,c} - \sum_{h=0}^{c-1} z_{i,h}$ for all $c \in \{0, 1, \dots, B\}$ and $i \in [n]$.

Proof: We provide a proof for completeness. Fix any agent $i \in [n]$. For the first property, we will show that $z_{i,c_1} \leq z_{i,c_2}$ for any values $c_1 < c_2$. By the truthfulness condition when i 's true valuation is c_1 and he bids c_2 ,

$$c_1 \cdot z_{i,c_1} - q_{i,c_1} \geq c_1 \cdot z_{i,c_2} - q_{i,c_2}.$$

Similarly when i 's true valuation is c_2 and he bids c_1 ,

$$c_2 \cdot z_{i,c_2} - q_{i,c_2} \geq c_2 \cdot z_{i,c_1} - q_{i,c_1}.$$

Adding the above two inequalities and rearranging, we get $(c_2 - c_1)(z_{i,c_2} - z_{i,c_1}) \geq 0$, i.e., $z_{i,c_2} \geq z_{i,c_1}$ as desired. This proves the monotonicity of $z_{i,*}$.

For the second property, fix also any value c . For each $h \leq c$, when i 's true valuation is h and he bids $h - 1$, by truthfulness:

$$q_{i,h} - q_{i,h-1} \leq h \cdot z_{i,h} - h \cdot z_{i,h-1}.$$

Adding this inequality over all $h \in \{1, \dots, c\}$,

$$q_{i,c} - q_{i,0} \leq c \cdot z_{i,c} - \sum_{h=0}^{c-1} z_{i,h}.$$

Now, voluntary participation implies that $q_{i,0} = 0$, which proves the desired inequality. \blacksquare

Also define $x_i := \sum_{c=0}^B \Pr[v_i = c] \cdot z_{i,c}$ for each $i \in [n]$. This denotes the probability that agent i is served by the mechanism when all agents bid their true valuation.

Claim 4.3 $\{x_i : i \in [n]\} \in \mathcal{P}(\mathcal{I})$.

Proof: The feasibility constraint imposed by \mathcal{I} implies that $\{X_i(\mathbf{b}) : i \in [n]\} \in \mathcal{I}$ for each $\mathbf{b} \in \{0, 1, \dots, B\}^n$. Since $\mathcal{P}(\mathcal{I})$ is a relaxation of \mathcal{I} , it is clear that $\{X_i(\mathbf{b}) : i \in [n]\} \in \mathcal{P}(\mathcal{I})$ for all \mathbf{b} .

Since $\mathcal{P}(\mathcal{I})$ is a convex set, we have $\sum_{\mathbf{b}} \lambda(\mathbf{b}) \cdot \{X_i(\mathbf{b})\} \in \mathcal{P}(\mathcal{I})$ for any convex multipliers λ s. Setting $\lambda(\mathbf{b}) := \Pr[v_j = b_j \forall j \in [n]] = \prod_{j=1}^n \Pr[v_j = b_j]$, for each $i \in [n]$ we have $\sum_{\mathbf{b}} \lambda(\mathbf{b}) \cdot X_i(\mathbf{b})$ equal to

$$= \sum_{c=0}^B \Pr[v_i = c] \cdot \sum_{\mathbf{b}_{-i}} \Pr[v_j = b_j : j \neq i] \cdot X_i(\mathbf{b}_{-i}, c) = \sum_{c=0}^B \Pr[v_i = c] \cdot z_{i,c} = x_i.$$

Thus we obtain $\{x_i\} \in \mathcal{P}(\mathcal{I})$ as claimed. \blacksquare

Combining Lemma 4.2 and Claim 4.3 we obtain the following LP relaxation for valid mechanisms:

$$LP_M = \max \sum_{i \in [n]} \sum_{c=0}^B \Pr[v_i = c] \cdot \left[c \cdot z_{i,c} - \sum_{h=0}^{c-1} z_{i,h} \right] \quad (8)$$

$$\text{subject to} \quad 0 \leq z_{i,0} \leq z_{i,1} \leq \dots \leq z_{i,B} \leq 1 \quad \forall i \in [n] \quad (9)$$

$$x_i = \sum_{c=0}^B \Pr[v_i = c] \cdot z_{i,c} \quad \forall i \in [n] \quad (10)$$

$$\{x_i : i \in [n]\} \in \mathcal{P}(\mathcal{I}). \quad (11)$$

We note that this LP is, in fact, a *relaxation* of the optimal mechanism, and the optimal value of this LP may be strictly larger than that of the optimal mechanism. (E.g. for a single matroid constraint, this gap can be as large as $\frac{e}{e-1}$.) It is also known that this gap can be closed by adding exponentially many valid inequalities—the so-called *Border inequalities*. However, this is not important for the current development, and the interested reader may refer to [24] for a thorough treatment of this area.

We are now ready to relate the above two LPs: LP_M , which is a relaxation of the optimal mechanism, and LP_P , our relaxation of the stochastic probing instance.

Lemma 4.4 $LP_P \geq LP_M$. Hence the optimal LP value of the stochastic probing instance is at least the revenue of the optimal mechanism.

Proof: Given any feasible solution $\langle z_{i,c}, x_i \rangle$ to LP_M , we construct a feasible solution $\langle x_{i,c}, y_{i,c} \rangle$ to LP_P of the same objective value. Set $y_{i,c} := z_{i,c} - z_{i,c-1}$ for all $i \in [n]$ and $c \in \{0, 1, \dots, B\}$; using $z_{i,-1} = 0$. Note that $\mathbf{y} \geq 0$ due to constraint (9). Also $\sum_{c=0}^B y_{i,c} = z_{i,B} \leq 1$ for each $i \in [n]$. This shows that constraints (6)-(7) in LP_P are satisfied.

Since $x_{i,c} = \Pr[v_i \geq c] \cdot y_{i,c} = \Pr[v_i \geq c] \cdot (z_{i,c} - z_{i,c-1})$, we have for each $i \in [n]$,

$$\sum_{c=0}^B x_{i,c} = \sum_{c=0}^B \Pr[v_i \geq c] \cdot (z_{i,c} - z_{i,c-1}) = \sum_{c=0}^B z_{i,c} \cdot (\Pr[v_i \geq c] - \Pr[v_i \geq c+1]),$$

which equals $\sum_{c=0}^B z_{i,c} \cdot \Pr[v_i = c] = x_i$. Thus constraint (11) in LP_M implies constraint (5) in LP_P . Finally, the objective value (3) of LP_P is:

$$\begin{aligned} \sum_{i \in [n]} \sum_{c=0}^B c \cdot x_{i,c} &= \sum_{i \in [n]} \sum_{c=0}^B c \cdot \Pr[v_i \geq c] \cdot (z_{i,c} - z_{i,c-1}) \\ &= \sum_{i \in [n]} \sum_{c=0}^B z_{i,c} \cdot (c \cdot \Pr[v_i \geq c] - (c+1) \cdot \Pr[v_i \geq c+1]) \end{aligned}$$

$$\begin{aligned}
&= \sum_{i \in [n]} \sum_{c=0}^B z_{i,c} \cdot (c \cdot \Pr[v_i = c] - \Pr[v_i \geq c+1]) \\
&= \sum_{i \in [n]} \sum_{c=0}^B \Pr[v_i = c] \cdot c \cdot z_{i,c} - \sum_{i \in [n]} \sum_{c=0}^B z_{i,c} \cdot \sum_{h=c+1}^B \Pr[v_i = h] \\
&= \sum_{i \in [n]} \sum_{c=0}^B \Pr[v_i = c] \cdot c \cdot z_{i,c} - \sum_{i \in [n]} \sum_{c=0}^B \Pr[v_i = c] \cdot \sum_{h=0}^{c-1} z_{i,h},
\end{aligned}$$

which is exactly the objective (8) of LP_M . ■

5 Unweighted Probing with Deadlines

In this section we consider a generalization of the stochastic probing problem in the presence of *global time*. Each probe requires one unit of time and each element $e \in V$ has a *deadline* d_e (in the global time) after which it expires. As before, we have inner \mathcal{I}_{in} and outer \mathcal{I}_{out} packing constraints on the set of chosen and probed elements respectively. We show that a natural greedy algorithm achieves a good approximation for *unit-weighted stochastic probing with deadlines*, when the inner and outer constraints are k -systems.

Theorem 5.1 *There is a $\frac{1}{2(k_{in}+k_{out}+1)}$ -approximation algorithm for unweighted stochastic probing with deadlines, when \mathcal{I}_{in} and \mathcal{I}_{out} are k_{in} - and k_{out} -systems.*

The main idea is to relax the global deadline constraints into an *outer laminar matroid* constraint \mathcal{L} , and then relate the deadline probing problem to the usual probing problem with outer constraints $\mathcal{I}_{out} \cap \mathcal{L}$ and inner constraints \mathcal{I}_{in} . The laminar matroid \mathcal{L} is defined as follows.

$$\mathcal{L} := \{U \subseteq V : |U \cap \{e : d_e \leq t\}| \leq t, \forall t \geq 1\}$$

Notice that the sets $D_t = \{e : d_e \leq t\}$ for $t \geq 1$ form a *chain family*⁴, and so \mathcal{L} is indeed a laminar matroid.

Algorithm 5.1 Greedy Algorithm for Unweighted Probing with Deadlines

```

1:  $Q \leftarrow \emptyset, S \leftarrow \emptyset, B \leftarrow \emptyset$  and  $t \leftarrow 1$ .
2: for  $e$  in non-increasing order of  $p_e$  value do
3:   if  $Q \cup \{e\} \in \mathcal{I}_{out} \cap \mathcal{L}$  then
4:     if  $S \cup \{e\} \in \mathcal{I}_{in}$  then
5:        $Q \leftarrow Q \cup \{e\}$  (i.e., potentially probe  $e$ )
6:       if  $t \leq d_e$  then
7:         probe element  $e$ , and  $t \leftarrow t + 1$ .
8:         if  $e$  active (happens with probability  $p_e$ ) then
9:            $S \leftarrow S \cup \{e\}$  (i.e., pick  $e$ )
10:      else
11:         $B \leftarrow B \cup \{e\}$ .
12:         $S \leftarrow S \cup \{e\}$  with probability  $p_e$ ; and  $S \leftarrow S$  otherwise.

```

The variable t in Algorithm 5.1 tracks the global time, which increases by one after each probe. The probed elements are $Q \setminus B$ and the chosen elements are $S \setminus B$. Observe that the algorithm defines a feasible policy since elements are only probed before their respective deadlines, and both inner \mathcal{I}_{in} and outer \mathcal{I}_{out} constraints are satisfied. We use the sets Q, S, B to couple:

⁴A chain family is a collection of subsets $D_1 \subseteq D_2 \subseteq \dots \subseteq D_n$.

- this algorithm for the deadline probing instance \mathcal{J} , and
- the greedy algorithm (Section 2) for the usual probing instance \mathcal{K} having inner constraints \mathcal{I}_{in} and outer constraints $\mathcal{I}_{out} \cap \mathcal{L}$.

Clearly, any feasible policy for \mathcal{J} is also feasible for \mathcal{K} . Note that the greedy algorithm for \mathcal{K} will probe elements Q and choose elements S . This is the reason why sets Q and S are updated even when a probe does not occur in \mathcal{J} . Also, B denotes the set of elements that are probed in \mathcal{K} but not in \mathcal{J} . Consider a decision path π down the decision trees associated with \mathcal{J} and \mathcal{K} ; note that we couple instantiations in the two decision trees. By the analysis in Section 2, the algorithm's objective in \mathcal{K} is

$$\mathbb{E}[\text{alg}(\mathcal{K})] = \sum_{\pi} \Pr(\pi) \cdot \sum_{e \in Q_{\pi}} p_e \geq \frac{\text{opt}(\mathcal{K})}{k_{in} + k_{out} + 1}.$$

Recall that \mathcal{K} has an inner k_{in} -system \mathcal{I}_{in} and outer $(k_{out} + 1)$ -system $\mathcal{I}_{out} \cap \mathcal{L}$. Moreover, the algorithm's objective in \mathcal{J} is

$$\mathbb{E}[\text{alg}(\mathcal{J})] = \sum_{\pi} \Pr(\pi) \cdot \sum_{e \in Q_{\pi} \setminus B_{\pi}} p_e.$$

The next lemma relates these two quantities.

Lemma 5.2 *For any outcome π , $\sum_{e \in Q_{\pi}} p_e \leq 2 \cdot \sum_{e \in Q_{\pi} \setminus B_{\pi}} p_e$.*

Proof: This proof also relies crucially on the greedy ordering in terms of probabilities. Note that each element $e \in B_{\pi}$ must have been considered at time $t > d_e$. Moreover, time t is increased only by elements $Q_{\pi} \setminus B_{\pi}$. Now by the greedy ordering,

$$|(Q_{\pi} \setminus B_{\pi}) \cap \{f : p_f \geq p_e\}| \geq d_e, \quad \forall e \in B_{\pi}. \quad (12)$$

Furthermore, since $Q_{\pi} \in \mathcal{L}$ we also have $B_{\pi} \subseteq Q_{\pi} \in \mathcal{L}$. So,

$$|B_{\pi} \cap \{f : d_f \leq d_e\}| \leq d_e, \quad \forall e \in B_{\pi}. \quad (13)$$

Consider a bipartite graph H with left vertices B_{π} and right vertices $Q_{\pi} \setminus B_{\pi}$, with an edge between $e \in B_{\pi}$ and $f \in Q_{\pi} \setminus B_{\pi}$ iff $p_e \leq p_f$. We claim that there is a left-saturating matching in H . It suffices to show Hall's condition that for any subset $R \subseteq B_{\pi}$, its neighborhood $|\Gamma(R)| \geq |R|$. Let $e := \arg \max\{d_g : g \in R\}$. Then, we have $|R| \leq d_e$ using (13), and $|\Gamma(R)| \geq |\Gamma(e)| \geq d_e$ by (12). Since graph H has a left-saturating matching, it is clear that $\sum_{e \in B_{\pi}} p_e \leq \sum_{e \in Q_{\pi} \setminus B_{\pi}} p_e$. ■

Using this lemma, and the above bounds for $\text{alg}(\mathcal{J})$ and $\text{alg}(\mathcal{K})$,

$$\text{alg}(\mathcal{J}) \geq \frac{\text{opt}(\mathcal{K})}{2(k_{in} + k_{out} + 1)} \geq \frac{\text{opt}(\mathcal{J})}{2(k_{in} + k_{out} + 1)}.$$

The last inequality uses the fact that instance \mathcal{K} is a relaxation of instance \mathcal{J} . This proves the first part of Theorem 5.1.

Stochastic matching with deadlines. We give an application of Theorem 5.1 in the kidney exchange setting. Consider a set of patients in a hospital, where each patient j is expected to be in the system for d_j days. For each pair i, j of patients, there is a probability $p_{i,j}$ of having a successful match. On each day, the hospital can perform one compatibility test and surgery between some pair of patients. If patient j is not matched by day d_j , he/she is assumed to have left the

system. The difference from the usual stochastic matching [12, 3] is that the “timeout level” of each patient decreases every day, irrespective of whether he is probed. The goal is to schedule tests so as to maximize the expected number of matched patients. This can be modeled as the probing problem with deadlines, on groundset V being the edges of the complete graph on patients. Each edge (i, j) has deadline $\min\{d_i, d_j\}$ and probability $p_{i,j}$. There is no outer constraint, and the inner constraint requires the chosen edges to form a matching (2-system). Hence Theorem 5.1 implies a $\frac{1}{6}$ -approximation algorithm for this problem.

We note that an LP based approach as in [3] can also be used to obtain an approximation ratio of $1/6$ for this problem. However the above greedy algorithm is much simpler and extends to general k -system constraints. For the weighted case, our result (Theorem 3.4) does not seem to extend directly to this setting of deadlines. We leave this as an open question.

Acknowledgments

We thank Shuchi Chawla, Bobby Kleinberg, Tim Roughgarden, Rakesh Vohra, and Matt Weinberg for helpful clarifications and discussions. We also thank an anonymous reviewer for pointing out that the LP for weighted stochastic probing can be solved approximately for general k -systems, which leads to an approximation algorithm for weighted probing under k -system constraints. Part of this work was done when the first-named author was visiting the IEOR Department at Columbia University, and IBM Thomas J. Watson Research Center; he thanks them for their generous hospitality.

References

- [1] Marek Adamczyk. Improved analysis of the greedy algorithm for stochastic matching. *Inf. Process. Lett.*, 111(15):731–737, 2011.
- [2] Aaron Archer and Éva Tardos. Truthful mechanisms for one-parameter agents. In *FOCS*, pages 482–491, 2001.
- [3] Nikhil Bansal, Anupam Gupta, Jian Li, Julián Mestre, Viswanath Nagarajan, and Atri Rudra. When LP Is the Cure for Your Matching Woes: Improved Bounds for Stochastic Matchings. *Algorithmica*, 63(4):733–762, 2012.
- [4] Nikhil Bansal, Nitish Korula, Viswanath Nagarajan, and Aravind Srinivasan. On k -column sparse packing programs. In *IPCO*, pages 369–382, 2010.
- [5] Anand Bhalgat. A $(2 + \epsilon)$ -approximation algorithm for the stochastic knapsack problem. *Manuscript*, 2011.
- [6] Anand Bhalgat, Ashish Goel, and Sanjeev Khanna. Improved approximation results for stochastic knapsack problems. In *SODA*, pages 1647–1665, 2011.
- [7] Sayan Bhattacharya, Gagan Goel, Sreenivas Gollapudi, and Kamesh Munagala. Budget constrained auctions with heterogeneous items. In *STOC*, pages 379–388, 2010.
- [8] Liad Blumrosen and Thomas Holenstein. Posted prices vs. negotiations: an asymptotic analysis. In *ACM Conference on Electronic Commerce*, 2008.
- [9] Shuchi Chawla, Jason D. Hartline, David L. Malec, and Balasubramanian Sivan. Multi-parameter mechanism design and sequential posted pricing. In *STOC*, pages 311–320, 2010.

- [10] Chandra Chekuri, Marcelo Mydlarz, and F. Bruce Shepherd. Multicommodity demand flow in a tree and packing integer programs. *ACM Transactions on Algorithms*, 3(3), 2007.
- [11] Chandra Chekuri, Jan Vondrák, and Rico Zenklusen. Submodular function maximization via the multilinear relaxation and contention resolution schemes. In *STOC (Full version: <http://arxiv.org/abs/1105.4593>)*, pages 783–792, 2011.
- [12] Ning Chen, Nicole Immorlica, Anna R. Karlin, Mohammad Mahdian, and Atri Rudra. Approximating Matches Made in Heaven. In *ICALP (1)*, pages 266–278, 2009.
- [13] Brian C. Dean, Michel X. Goemans, and Jan Vondrák. Adaptivity and approximation for stochastic packing problems. In *SODA*, pages 395–404, 2005.
- [14] Brian C. Dean, Michel X. Goemans, and Jan Vondrák. Approximating the stochastic knapsack problem: the benefit of adaptivity. *Math. Oper. Res.*, 33(4):945–964, 2008.
- [15] M.L. Fisher, G.L. Nemhauser, and L.A. Wolsey. An analysis of approximations for maximizing submodular set functions II. *Mathematical Programming Study*, 8:73–87, 1978.
- [16] Sudipto Guha and Kamesh Munagala. Approximation algorithms for budgeted learning problems. In *STOC*, pages 104–113. 2007.
- [17] Klaus Jansen. Approximate strong separation with application in fractional graph coloring and preemptive scheduling. *Theor. Comput. Sci.*, 302(1-3):239–256, 2003.
- [18] T.A. Jenkyns. The efficiency of the “greedy” algorithm. In *7th South Eastern Conference on Combinatorics, Graph Theory and Computing*, pages 341–350, 1976.
- [19] Robert Kleinberg and S. Matthew Weinberg. Matroid prophet inequalities. In *STOC*, pages 123–136, 2012.
- [20] Rolf H. Möhring, Andreas S. Schulz, and Marc Uetz. Approximation in stochastic scheduling: the power of LP-based priority policies. *Journal of the ACM (JACM)*, 46(6):924–942, 1999.
- [21] R. Myerson. Optimal auction design. *Mathematics of Operations Research*, 6:58–73, 1981.
- [22] Tuomas Sandholm and Andrew Gilpin. Sequences of take-it-or-leave-it offers: near-optimal auctions without full valuation revelation. In *5th International Joint Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1127–1134, 2006.
- [23] A. Schrijver. *Combinatorial Optimization*. Springer, 2003.
- [24] R. Vohra. *Mechanism Design: A Linear Programming Approach*. Cambridge University Press, 2011.
- [25] Qiqi Yan. Mechanism Design via Correlation Gap. In *SODA*, pages 710–719, 2011.

A Bad Examples for Simpler LP-Rounding Algorithms

Here we observe that some natural LP-rounding algorithms that work for stochastic matchings [3] do not work in the setting of general matroids. Let (x, y) denote a solution to the linear relaxation \mathcal{LP} . Consider rounding this solution by considering elements to probe in the following order, where each element e is probed with probability $b \cdot y_e$ when permitted by the inner and outer constraints ($0 < b \leq 1$ is some constant).

- *Decreasing w_e value.* There is no inner constraint, and the outer constraint is a graphic matroid on the graph G (see Figure 1) consisting of edges $E := \{e_i\}_{i=1}^n \cup \{f_i\}_{i=1}^n$ and g . The weights on edges E are $M \gg 1$ each, and $w(g) = 1$. The probabilities on edges E are $\epsilon \ll \frac{1}{nM}$ each, and $p(g) = 1$. The fractional solution y has value one on edge g and value $1/2$ on each of E ; the LP objective is at least one. The expected weight from E is at most $2nM\epsilon$. Since edge g appears last in this order, the probability that g is *not* blocked by the outer graphic matroid is at most $(1 - b^2/4)^n$. (Note that if any edge of E is blocked then so is g .) So the expected total weight of the rounding algorithm is at most $2nM\epsilon + (1 - b^2/4)^n \ll 1$.
- *Decreasing p_e value.* Again, there is no inner constraint and the outer constraint is a graphic matroid on G (see Figure 1). The weights on E are one, and $w(g) = L \gg n$. The probabilities on E are one, and $p(g) = 1/2$. y has value one on edge g and value $1/2$ on each of E ; so the LP objective is at least $L/2$. The expected weight from E is at most $2n \ll L$. As before, since edge g appears last in this order, the expected weight from g is at most $L \cdot (1 - b^2/4)^n \ll L$. Hence the expected total weight is $\ll L$.
- *Decreasing $w_e \cdot p_e$ value.* There is no outer constraint and the inner constraint is a graphic matroid on graph H (see Figure 1), which consists of edges $E := \{e_i\}_{i=1}^n \cup \{f_i\}_{i=1}^n$ and $E' := \{g_j\}_{j=1}^N$. We set $N = n^2$. The weights are two on E , and N on E' . The probabilities are $1/3$ on E , and $\frac{1}{3N}$ on E' . y has value one on all edges, and the LP objective is at least $N/3$. This order puts edges of E before edges of E' . The probability that any particular edge of E' is *not* blocked in the inner graphic matroid is at most $(1 - b^2/9)^n \ll 1$. Thus the expected weight from E' is at most $N^2 \cdot (1 - b^2/9)^n \cdot \frac{1}{3N} \ll N$. The expected weight from E is at most $4n \ll N$. So the expected total weight is again much lesser than the LP objective.




Figure 1: Graphic matroids on G and H .